

How Well Are Two Process Models Standing Up to the Bayesian Challenge?

Steven Sloman
Cognitive & Linguistic Sciences
Brown University

What I learned in Cambridge this summer

Steven Sloman
Cognitive & Linguistic Sciences
Brown University

1. Cambridge is a lovely town.
2. You can only hear so many talks about two systems of thought in one day.
3. What I think of the “System 1” and “System 2” language
 - Allows talk about systems without committing to too much
 - Allows talk about systems without committing to anything
4. This 2 systems idea has been wildly overinterpreted.
 - The domain of application desperately needs to be contained: reasoning, not craving, walking, or even language.
 - Theories that bite off too much don’t say anything at all; must stay close to data.
 - I thought at first even Jonathan went too far (reasoning, decision-making, social cognition, and learning). E.g., learning and inference are fundamentally different issues (different time scales)
 - However, one of my favorite ideas was Jonathan’s: “one system trains the other.”

5. In some cases, domain of application too narrow: hypothetical thought doesn't strike me as a natural kind.
6. A special license should be required to draw boxes in Powerpoint.
 - What's the value of architectural distinctions?
 - I think they are a fabulous way of describing functional relations amongst theoretical entities. But they don't stand on their own (as evident in many talks, e.g., Reyna etc.)
 - Not mere neuroscience envy. We really saw the value of neuroimaging data for the enterprise of distinguishing systems.
 - o Kahneman: building an archive to turn psychology into a cumulative science.

7. Importance of computational models.

- Frankish: personal vs. subpersonal
- Smolensky: symbolic vs. subsymbolic

Computational distinction:

System 2: compositional rule following (Goel)

System 1: statistical structure.

➔ We'll never achieve any consensus on how the systems work, how they interact, how to describe them formally, and how they are implemented in the brain without agreeing on what they are trying to accomplish.

The Descriptive Bayesian Challenge

- The mind is always doing one thing, so it should always do whatever it's doing in the same way.
- People have a unique goal when reasoning: to choose actions guided by knowledge of how probable hypotheses are given the data. What we do is incorporate the likelihood of facts about a case given relevant hypotheses with prior beliefs about the distribution of those hypotheses via Bayes' rule:

$$P(H | D) = \frac{P(D|H)P(H)}{P(D)}$$

- a. beliefs are all probabilistic → no such thing as certain conclusions. Unlike TM, dual systems not required.
- b. belief formation considers both diagnosticity of data and prior belief → how we do so is the real problem

A General Answer

- Most problems have no verifiably optimal solution, so we have to approximate.
- Different approximations for different situations.
 - E.g., chess
 - Some search
 - Some pattern recognition

Dual Process Computational Theories That I Know About

- Sloman: statistical description of the environment vs. description in terms of production rules enforcing hard constraints.
- Evans & Over: pursuit of individual goals vs. normative prescriptions.
- Stanovich & West: genes' vs. individual's goals.

Do the theories meet the challenge?

- Yes. Merely assimilating phenomena to an equation fails to explain any of the dissociations we know about (working memory load, time pressure, Criterion S, NEXT SLIDE etc.)
 - Even if it's true that beliefs are all probabilistic, there may be (in fact, must be) multiple ways to arrive at them. The constraints imposed by the Bayesian challenge are easily met by many processing systems be they associative or rule-based.
- ➔ Nevertheless, the right sort of dual-process theory should spell out both computational constraints and say something about the underlying process.
- System 2 may be about constructing or choosing the right abstract conceptual structure to use as a frame for reasoning.

Applying the process dissociation procedure to judgment under uncertainty

Mário B. Ferreira

Leonel Garcia-Marques

Steven J. Sherman

JPSP in press

Inductive reasoning and the process dissociation procedure

Jacoby (1991) introduced a procedure to estimate the contribution of two separable processes by placing them

- in opposition (the *exclusion condition -EC-*)
- in concert (the *Inclusion condition -IC-*)

EXCLUSION PROBLEM

Several psychologists interviewed a group of people. The group included 30 engineers and 70 lawyers. The psychologists prepared a brief summary of their impression of each interviewee. The following description was drawn randomly from the set of descriptions: Dan is 45. He is conservative, careful and ambitious. He shows no interest in political issues and spends most of his free time on his many hobbies, which include carpentry, sailing, and mathematical puzzles.

Which of the following is more likely?

- a) Dan is an engineer*
- b) Dan is a lawyer*

INCLUSION PROBLEM

Several psychologists interviewed a group of people. The group included 70 engineers and 30 lawyers. The psychologists prepared a brief summary of their impression of each interviewee. The following description was drawn randomly from the set of descriptions: Dan is 45. He is conservative, careful and ambitious. He shows no interest in political issues and spends most of his free time on his many hobbies, which include carpentry, sailing, and mathematical puzzles.

Which of the following is more likely?

- a) Dan is an engineer*
- b) Dan is a lawyer*

Estimates of associative and rule-based reasoning assuming independence

$$CR_{IC} = r + a - ra$$

$$IR_{EC} = a(1-r) = a - ra$$


CR = *Correct* responses

a=Associative
reasoning

IR = *Incorrect* responses

r=Rule-based
reasoning

RESULTS



Mean proportion estimates of rule-based and heuristic reasoning as a function of high and low cognitive load memory based-judgments.

Experiment	Manipulation	r		a	
1	Instructions: intuition vs. rationality	0.14 p< 0.05	0.39	0.68 n.s	0.74
2	Cognitive load: low vs. high	0.35 p< 0.05	0.21	0.77 n.s	0.76
3	Associative priming	0.32 n.s	0.27	0.70 p<0.05	0.82
4	Rule-based priming	0.33 p<0.05	0.45	0.81 n.s	0.80

The Normative Bayesian Challenge

Challenge:

- One kind of normative model is consistent with both fallacies and correct responses.
- Everything comes down to incorporating the likelihood of facts about a case given relevant hypotheses with prior beliefs about the distribution of those hypotheses.
- Given the right model of likelihood and the right priors, model is consistent with many patterns of performance.

The Normative Bayesian Challenge

e.g., Tenenbaum & Griffiths: The rational basis of representativeness:

- Explain conjunction fallacy in terms of selection of relevant hypotheses. Representativeness is determined not just by $P(\text{data} \mid \text{target hypothesis})$ but also $P(\text{data} \mid \text{alternative hypotheses})$.

HHHHH is less representative than HTTHTT because

$P(\text{HTTHTT} \mid \text{2-headed coin}) < P(\text{HHHHH} \mid \text{2-headed coin})$.

- They don't explain Linda at all.

Three Dual-process Responses

1. There exists one kind of rationality determined by the right choice of normative model. That choice is far from trivial, and it should be blind to participants' responses.
 - Note that rational is not the same as adaptive.
2. Depends on response, not on reasoning system. Errors generally more likely from associative system.

Three Dual-process Responses

2. Evans & Over (1996):

Rationality₁: Thinking, speaking, reasoning, making a decision, or acting in a way that is generally reliable and efficient for achieving one's goals.

- Not necessarily Bayesian

Rationality₂: Thinking, speaking, reasoning, making a decision, or acting when one has a reason for what one does sanctioned by a normative theory.

- Necessarily Bayesian

But: If behavior is “generally reliable and efficient for achieving a goal,” then that behavior must be doing something consistent with some hypothesis. We can always stipulate that hypothesis is the desired goal, in which case people are always **Rational₁**.

Three Dual-process Responses

3. Stanovich & West (2003): evolutionary vs. instrumental rationality:
 - Associative system pursues gene's interests (short-leash goals); rule-based system pursues individual interests (long-leash goals).
 - As long as this is a principled distinction (and I think it is), then it's irrelevant that the different computational goals can be described with the same mathematical model. The model must be systematically realized in different ways to account for the systematically different goals pursued by the two systems.
- ➔ The normative challenge is not a serious one.

A more problematic Bayesian view: Causal Reasoning

- A good chunk of human reasoning concerns causal relations among events
- Causality obeys a special kind of logic, a logic that conforms to the central role that intervention plays in reasoning about events in general, and about the effects of human action in particular.
- Causal reasoning is “Bayesian” in the sense that it has been modeled using Bayesian-belief systems
 - Causal Bayes nets are popular representations of causal beliefs that provide a host of conceptualizations about how causal beliefs are learned and how causal inferences are made.

Causal models are plausible psychological representations of uncertainty


- Formally, Bayes' nets are representations of probability distributions.
- Knowledge about likelihood and about how likelihoods change when we condition on events is encapsulated in a qualitative way by our causal beliefs. For example,

Inductive Inference

Sergey Blok

Garbage has toxin X.

Therefore, rats have toxin X.




Rats have toxin X.

Therefore, garbage has toxin X.

==> Both (observational) arguments strong.

Interventional case

Garbage is injected with toxin X.
Therefore, rats have toxin X.




Rats are injected with toxin X.
Therefore, garbage has toxin X.



==> Only first argument strong.

Result



The problem

- It's very difficult to make a principled distinction between associative and deliberative causal reasoning:
 - Causal reasoning isn't easily described as either associative or rule-based.

Both kinds of reasoning may be described using the same causal principles

- It's obvious that the logic of causality applies to deliberative causal reasoning. Most of the direct evidence for causal reasoning and for the logic of intervention in the domain of causal reasoning comes from the deliberative domain.
 - troubleshooting (what kind of stupid automated telephone system calls you back when you hang up on it?; causal syllogisms (A causes B, B prevents C, what does A do to C?))
- There's compelling evidence both that associative causal reasoning obeys the logic of causation and evidence that it doesn't.

Evidence for associative causal reasoning

- Pronoun resolution:
 - Steven admires Jonathan because he is so candid.
 - Steven annoys Jonathan because he is so candid.
- Glymour example: Having yellow teeth is correlated with lung cancer. Therefore, you should whiten your teeth to lower the probability of getting lung cancer. True or false?
- Experts troubleshooting complex causal systems. E.g., electricians.

Evidence against associative causal reasoning

- Medical decision-making literature: Only inexperienced students use causal models.
- Large implicit learning literature indicating correlations can be learned between arbitrary variables. In most such cases, there's no natural direction to the relation and no other reason to believe that a directional relation is being learned.
- Self-deception. Reasoning in such a case violates causal structure and must be unconscious to be useful.

Self-deception: Quattrone & Tversky

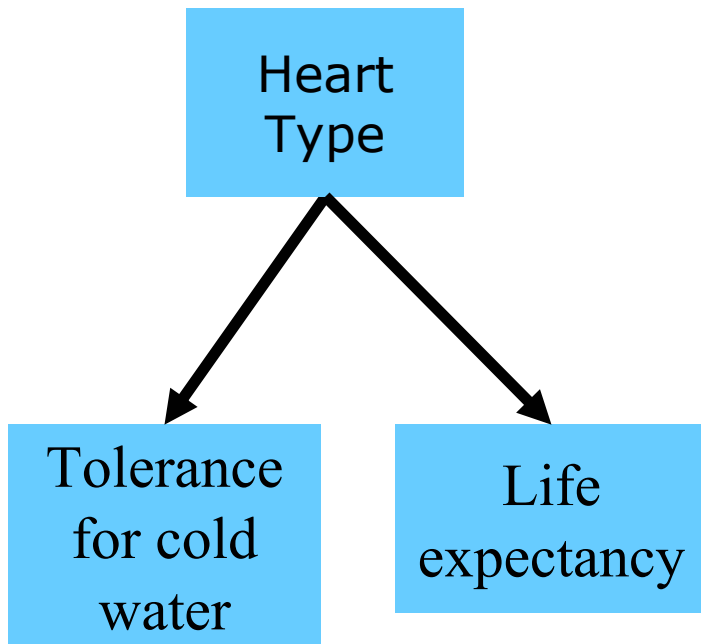
To Stanford undergraduates:

- “We’re studying how rapid changes in temperature affect heart rate after exercise.”
- Students asked to hold their arms in very cold water for as long as they could. Then asked to do it again after spending a minute vigorously riding an exercycle.
- Manipulated hypothesis presented:
 - Half told: people can tolerate cold water for longer after exercise if they have a good type of heart.
 - Other half the opposite: people can tolerate cold water for less time if they have the good heart type.

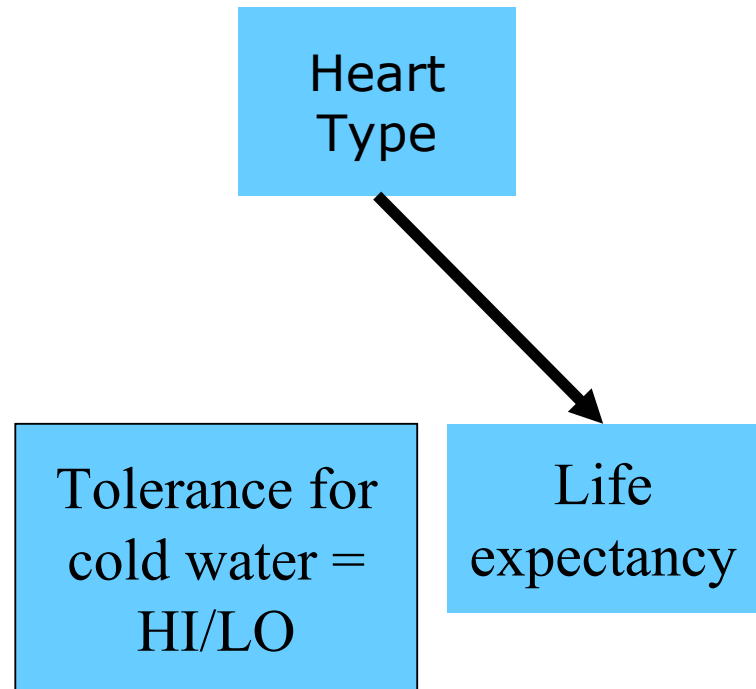
Result: group told that good heart = longer in fact lasted longer.

Causal Analysis

Pre-choice Model:



Model for choice:



Causality and Two Minds

- The million-dollar (540,862 GBP) question is “does causal associative reasoning obey the logic of causality?”
 - If not, we can rest assured that dual-process models are on the right track.
 - If so, then I’d have to concede that there’s a fairly well-understood representational theory that explains a good chunk of human reasoning, and it explains all types.
- System 2 required to generate novel causal structure and to administer interventions.
 - A la Stanovich’s reflective system but
 - i. After administration, reasoning could still be associative.
 - ii. Not limited to hypothetical thought; also counterfactual thought, prediction, causal attribution, etc.
 - iii. Something like the Ramsey test must be applicable even to counterfactual conditionals.

How Many Reasoning Systems Are There?

- The right question is: “How many **parallel** reasoning systems are there”?
- The answer is: 2.